AMENDMENTS TO THE CLAIMS

1-48. (Canceled)

49. (Previously presented) A method for searching a corpus of documents containing terms, comprising:

defining a knowledge domain;

identifying a set of reference documents in the corpus pertinent to the domain;

finding lexical affinities of the terms in the reference documents, the lexical affinities of a given term comprising other terms that co-occur with the given term in sentences in the reference documents, such that the other terms are separated from the given term in the sentences by no more than a predetermined number of words;

inputting a search query comprising query terms;

refining the search query using the lexical affinities of the query terms that were found in the reference documents; and

searching the corpus to find information in the domain responsive to the refined query.

50. (Previously presented) The method according to claim 49, wherein the search query comprises search terms, and wherein refining the search query comprises adding to the search terms further terms found to have lexical affinity to the search terms.

51-57. (Canceled)

58. (Previously presented) Apparatus for searching a corpus of documents containing terms, comprising:

a memory, adapted to store an identification of a set of reference documents in the corpus pertinent to a predefined knowledge domain; and

a search processor, which is adapted to find lexical affinities of the terms in the reference documents, the lexical affinities of a given term comprising other terms that co-occur with the given term in sentences in the reference documents, such that the other terms are separated from the given term in the sentences by no more than a predetermined number of words, and responsive to receiving a query comprising query terms as input, is adapted to refine the search query using the lexical affinities of the query terms that were found in the reference documents and to search the corpus to find information in the domain responsive to the refined query.

59-61. (Canceled)

62. (Previously presented) A computer software product for searching a corpus of documents, the product comprising:

a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to receive a definition of a knowledge domain and an identification of a set of reference documents in the corpus pertinent to the domain and to find lexical affinities of the terms in the reference documents, the lexical affinities of a given term comprising other terms that co-occur with the given term in sentences in the reference documents, such that the other terms are separated from the given term in the sentences by no more than a predetermined number of words, and further cause the computer, responsive to a query comprising query terms, to refine the search query using the lexical affinities of the query terms that were found in the reference documents and to search the corpus to find information in the domain responsive to the refined query.